SEPTEMBER 1994 ENTERPRISE SYSTEMS JOURNAL VOLUME IX • NUMBER 9

Focusing On IBM Host-Based Enterprise-Wide Computing

ENTERPRISE SYSTEMS JOURNAL

September 1994

SPRINT TRACKS CHANGES

Even The Difference In A Pin Drop

CONTENTS

Focusing On IBM Host-Based Enterprise-Wide Computing
ENTERPRISE SYSTEMS
J · O · U · R · N · A · L

September 1994, Volume 9, Number 9

When a problem occurs, one of the first questions asked is, "What has changed?" With 19 individual LPARs on differing mainframes distributed throughout the country, Sprint Corp. was faced with a change-tracking challenge of epic proportions. To find out how it met that challenge, turn to Philip E. Courtney's article on page 20. Cover illustration by Robert Kimmerle.

I/S Management

Operating System Platforms

Client/Server Enterprise Solutions

Database Management

Electronic Document Management

Contingency Planning

First Impression

Departments

ESJ Supplement

Understanding Version 2 Of DB2/6000 And DB2/2

By C. M. Saracco and Charles J. Bontempo

Object support, connectivity enhancements, parallel I/O processing and new optimization techniques are included in Version 2 of IBM's DB2 products for AIX/6000 and OS/2 platforms. Such features are designed to improve the competitiveness of these products by providing greater flexibility, improved access to other data sources and better overall performance.

Any discussion of such features is bound to raise questions. For example, what technology is IBM actually providing in these areas? What value might these technologies bring to users? These questions are best answered by describing each of these technologies and its potential benefits.

Object Extensions

Like many vendors, IBM has been working on incorporating a variety of object-oriented features into its relational DBMS products. As shown in Figure 1, such features include a more flexible type system, capable of accommodating multimedia data types as well as user-defined data types; additional integrity mechanisms (including triggers and constraints) that help make the DBMS more active and enable users to associate more meaning with their data; and support for database access through C++ and Smalltalk. These features and others are often considered "object" extensions because they provide some of the capabilities associated with object technology.

For example, relational DBMSes are sometimes criticized for supporting only "record-oriented" data (mainly, formatted fields of character strings and numeric data types), while object-oriented programming languages enable users to create "nontraditional" types, some of which may employ an inherently complex internal structure. How are vendors such as IBM addressing this?

DB2/6000 and DB2/2 Version 2 use 2GB large objects (sometimes called "BLOBs") to support a variety of nontraditional data types such as audio, compressed video, image and documents. A single table can contain a number of such large-object columns, and users can use various SQL functions and operators to access data within these columns.

Support for user-defined types further enhances the type system of DB2/6000 and DB2/2 Version 2. User-defined types enable users to define new types based on system-supplied types. In doing so, they can specify certain operational restrictions for these new types to capture the behavior applicable for these types. For example, an "account number" type can be derived from the DBMS' native support for integers. In doing so, users can specify that account numbers cannot be subjected to arithmetic operations, since adding or multiplying two account numbers would be considered invalid. This helps users capture more of the semantics they inherently associate with their data and can help prevent erroneous or misleading results.

A related object extension involves support for user-defined functions, which enables users to define new SQL scalar

DB2/6000 and DB2/2 Version2 enable users to control the number and types of techniques considered by the optimizer when determining a reasonable access path to the data.

FIGURE 1

Overview Of IBM's Object Extensions

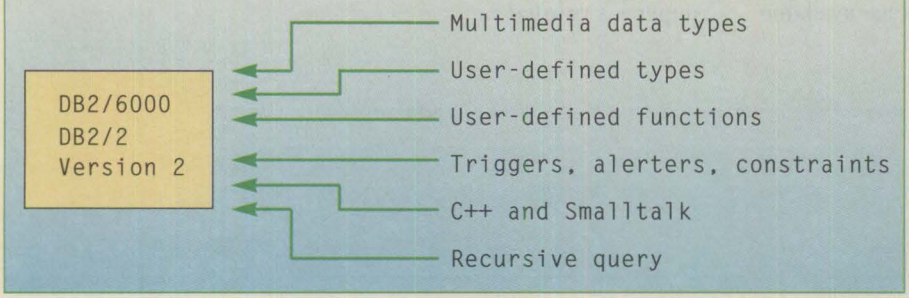

functions that can be incorporated into standard queries involving both system-supplied data types and user-defined data types. These functions, coded in a language such as C and registered with the DBMS through a CREATE FUNCTION statement, extend the breadth of operations supported by the DBMS. If desired, user-defined functions can be written to understand the internal structure of a user-defined type based on large objects. Such functions can work with various portions of this structure as desired.

Among the features of Version 2 that help provide for a more "active" environment and a broader range of integrity mechanisms are support for triggers, event alerters and constraints. Triggers enable users to instruct DB2/6000 or DB2/2 to automatically enforce various business policies whenever a given database event occurs (such as an attempt to insert or delete a row from a table, or update a given column of a table). Such activities can cause a trigger to execute, which, in turn, may cause some user-specified event to occur. This might include consulting data in other tables, updating other tables or rolling back the work.

IBM's trigger mechanism can also function as an event alerter, enabling a database event to cause a non-DBMS activity to automatically occur. This might include consulting or modifying a file, or sending an electronic message whenever someone makes a specified change to a named table.

Constraints represent another integrity mechanism incorporated into DB2/6000 and DB2/2 Version 2. They are useful for ensuring that the data values of one or more columns of a table fall within a certain range. For example, a constraint might ensure that serial numbers for all parts range from 1 to 1,000,000 or all salaried employees at the New York City office who work in department XYZ receive at least a $1000 yearly bonus.

Other object enhancements of DB2/6000 and DB2/2 Version 2 include support for recursive queries and two object-oriented programming languages. The latter involves support for embedded SQL in C++ as well as support for Smalltalk via VisualAge, an IBM tool that supports rapid application development.

Connectivity Enhancements

In addition to its first stage of object features, IBM also announced connectivity enhancements for its DB2/6000 and DB2/2 Version 2 products. These enhancements help improve access to both "live" and copied data.

DBMS enhancements account for the improved access to "live" data. In this area, IBM has enabled its relational DBMSes to support distributed unit of work processing in a LAN-based environment. This means a single transaction could update multiple DB2/6000 or DB2/2 databases, with the DBMS automatically coordinating the two-phase commit processing required to ensure the integrity of the transaction. If desired, a transaction manager (such as CICS/6000) can be used instead to coordinate the two-phase commit processing.

In addition, DB2/6000 and DB2/2 Version 2 can service requests from other DBMSes that employ the Distributed Relational Database Architecture (DRDA). A number of IBM and non-IBM DBMSes have announced or shipped support for DRDA "requestors," including DB2 (for MVS), DB2/400 (for AS/400), DB2/VM (formerly SQL/DS), Informix (Informix Software, Inc., Menlo Park, CA) and Oracle (Oracle Corp., Redwood Shores, CA). In previous releases, DB2/6000 and DB2/2 could only make requests of DRDA-enabled servers but could not respond to requests from DRDA-enabled products.

While these distributed database enhancements for "live" data access improve the connectivity of these DBMSes, a number of users have identified a need to maintain local copies of data that may be slightly out of sync with information in their (remote) production databases. This capability, which is sometimes referred to as "replication," can offer certain performance gains and provide an effective alternative for managing local and remote resources. This is possible because network traffic can be minimized or more readily controlled (particularly if incremental changes to copies are scheduled to be sent or received only at specific intervals) and dependencies on remote resources (including network and DBMS software) can be kept to a minimum.

In this area, IBM offers products that enable users to maintain copies of data stored in DB2 (on MVS and OS/400 platforms) in their local DB2/6000 and DB2/2 databases. In addition, data from certain nonrelational sources (such as IMS) can be copied into DB2/6000 and DB2/2 tables as well.

Because decision support applications are particularly well-suited to a replicated environment, IBM's Data Propagator Relational product enables users to manipulate the data as desired before the changes are propagated. This enables users to copy only portions of tables, joined tables or the results of aggregate functions applied to one or more tables. In addition, copies can be scheduled to occur at specific intervals (such as every hour or every day), and users may opt to have only the changes made since the last copy interval propagated.

Parallelism And Performance Issues

As with its other relational DBMSes, IBM has incorporated support for parallelism into its DB2 products for AIX and OS/2. One aspect of this involves support for parallel I/O processing. This enables multiple I/O processes to be executed simultaneously, so data is moved more rapidly from disks to database buffers to help improve response time. In addition, DB2/6000 and DB2/2 Version 2 can perform parallel I/O processing for backup and recovery tasks, which can also improve response time.

A related product, DB2 Parallel Edition, was designed to exploit IBM's POWERparallel System hardware and support read/write activities in a multiprocessor environment. This multiprocessor environment is based on RISC System/6000 technology and involves multiple nodes connected via a LAN or high-speed switch in a shared-nothing environment. Each node runs its own copy of the AIX operating system and the DBMS, and each has its own memory and disk(s). Queries can be decomposed into separate subtasks, with each node executing these tasks in parallel on its own portion of the data.

Note that the hardware architecture of this approach differs from that associated with IBM's System/390 Parallel Query Server, an MVS-based product supported by DB2. The S/390 Parallel Sysplex (on which the Parallel Query Server runs) enables multiple nodes to share data on one or more disks through the use of high-speed fiber optic channels (ESCON). Each node itself consists of multiple processors configured in a shared-memory environment.

Optimization Enhancements

Although support for various forms of parallelism can certainly offer users performance gains, other DBMS issues must be considered as well. For example, query optimization can be critical for achieving high performance in a relational DBMS environment. This was one area of focus for DB2/6000 and DB2/2 Version 2, which feature many optimization enhancements, some of which were drawn from IBM's STARBURST research project.

One optimization improvement enables the DBMS to automatically transform many complex queries into more efficient forms, helping to ensure that good performance is achieved regardless of how the query is expressed. In addition, the optimizer has been enhanced to consider a greater number of alternatives and more accurately estimate the cost of these. Because considering more alternatives can involve increased resource usage, DB2/6000 and DB2/2 Version 2 enable users to control the number and types of techniques considered by the optimizer when determining a reasonable access path to the data. This enables users to adjust the search space of optimization or specify the level of optimization desired. By using this feature, users can fine tune their sys-

tems based on the particular queries issued by an application and the resources available on a particular platform.

Another tuning mechanism commonly used is the EXPLAIN facility, which can help programmers and administrators understand the access plans selected by the optimizer for various queries. Understanding the access plan selected can help users tune the DBMS, perhaps by adding an index on a table or updating catalog statistics. DB2/6000 and DB2/2 Version 2 provide more detailed information through the EXPLAIN facility — information that is stored in tables to provide easy access. This EXPLAIN statement will support the syntax of the DB2 (MVS) EXPLAIN statement.

In addition, Version 2 enables users to update certain catalog statistics that influence optimization, changing these statistics and reviewing subsequent access plan information to perform "what if" analysis on their queries. This capability also enables users working in a test environment to enter statistics they believe would be more characteristic of a production environment, thereby giving them a chance to make some early performance estimates and experiment with various alternatives.

Other performance enhancements include a high-speed LOAD utility, additional database monitoring capabilities, additional support for prefetching data and support for tablespaces. The latter feature also offers administrators more control over the placement of their data as well as the ability to back up and restore their data at a more granular level (at a tablespace level).

ABOUT THE AUTHOR

C.M. Saracco has 10 years experience working with both object and relational databases. Together with Charles Bontempo, she is currently writing a book on database topics to be published by Addison-Wesley. She works at IBM Software Solutions Division, P.O. Box 49023, Office C360, San Jose CA 95161, (408) 463-3107.

Charles J. Bontempo has more than 30 years experience in the database arena. He teaches and consults on database topics for IBM and at the Polytechnic University in New York. IBM Software Solutions Division, Route 100, Office 4D-06, Somers, NY 10589, (914) 766-1394.